

For release to Computing Research News
February 2002
Pamela P. Walatka
pwalatka@mail.arc.nasa.gov
1646 words

Adventures in Computational Grids

Sometimes one supercomputer is not enough. Or your local supercomputers are busy, or not configured for your job. Or you don't have any supercomputers. You might be trying to simulate worldwide weather changes in real time, requiring more compute power than you could get from any one machine. Or you might be collecting microbiological samples on an island, and need to examine them with a special microscope located on the other side of the continent. These are the times when you need a computational grid.

A computational grid is a network of geographically and organizationally distributed resources such as computers, instruments, and data, giving a user single-sign-on access to all the resources on the grid. These resources are

managed by diverse organizations in widespread locations, and shared by researchers from many different institutions. Grid users establish their identities by getting authentication certificates, obtain accounts on the grid's computational resources, and then use the resources—which are often scattered throughout a continent or beyond—from their desktops.

The idea of computational grids took shape in a book edited by Ian Foster and Carl Kesselman, The Grid: Blueprint for a Future Computing Infrastructure, Morgan-Kaufmann, 1999. Foster and Kesselman initiated a project, Globus, to develop enabling software for grids: infrastructure, services, and application toolkits; Steve Tuecke, Lee Liming and dozens of developers are responsible for continued improvement. (See www.globus.org for more information.) The ambitious Globus project provides free open source software for a number of developing grids. Other organizations are developing alternative software, but none are as widely used as the Globus toolkit.

FROM DREAMWARE TO TESTBEDS

Computational grids have progressed from dreamware, through trials (and tribulations), to the existence of persistent testbeds, with production capabilities.

Universities without supercomputers have strung together clusters of workstations to provide their researchers with high performance resources. Businesses have begun to explore the possibilities in grid computing; IBM is involved (see IBM: Linux To Seed Next Generation Of Grid Computing. CRN, January 30, 2002, and <http://www.globus.org/about/news/HPCwire/ibmGridsTransform12-7-01.html>). Irving Wladawsky-Berger, Vice President, Technology and Strategy, IBM Server Group, calls grid computing the "key to advancing e-business into the future and the next step in the evolution of the Internet towards a true computing platform." He predicts that grid computing, like supercomputing before it, provide a vast infrastructure for e-business. (<http://www-1.ibm.com/servers/events/grid.html>) Great Britain is developing a grid, and Euroglobus is underway (see <http://www.euroglobus.unile.it/>). The Global Grid Forum (<http://www.gridforum.org/>) meets semiannually to consolidate standards for grids.

Eventually there may be one global grid, similar to the web but giving users access to high performance resources, with the ability to actually run jobs on those resources, not just interact with what is being served. On the road to a single worldwide grid, testbed grids are being developed.

The two most substantial grids in terms of compute cycles already provided to grid users, persistence, and scope of resources, are:

- The NSF PACI National Technology Grid
- The NASA Information Power Grid.

The National Science Foundation PACI Program

(<http://www.interact.nsf.gov/cise/descriptions.nsf/pd/paci?OpenDocument>)

combines the National Computational Science Alliance (NCSA) in Urbana-Champaign, Illinois, and the National Partnership for Advanced Computational Infrastructure (NPACI), at the San Diego Supercomputer Center in San Diego, California. This PACI National Technology Grid provides interconnected high-performance computing systems and powerful instruments, and has pioneered the development, application, and testing of grid infrastructure. For example, Randy Butler's group at NCSA provides the Grid-in-a-Box

(<http://www.ncsa.uiuc.edu/TechFocus/Projects/NCSA/Grid-in-a-Box.html>),

which provides the software and documentation to enable users to rapidly build and deploy functional grids. Also at NCSA, Doru Marcusiu manages collaborative grid infrastructure development with NASA's IPG and the San

Diego part of NPACI. In San Diego, Mary Thomas, Keith Thompson, Steve Mock, and many others have developed a Hot Page for interactive access to their grid.

NASA created the Information Power Grid (IPG)
(<http://www.ipg.nasa.gov/>). IPG makes available supercomputers, high-end scientific instruments, and terabyte datasets. A web portal, Launchpad, (<http://www.ipg.nasa.gov/launchpad/servlet/launchpad>) provides IPG users the ability to submit jobs to "batch" compute engines located at NASA centers across the United States, execute commands on these resources, transfer files between two systems, obtain status on systems and jobs, and modify the user's environment

The IPG team has pioneered grid development in the areas of automated parameter studies, grid services, system status, data mining, Globus security, performance monitoring, benchmarks, documentation, system availability, and testing. Issac Lopez heads the IPG team at NASA Glenn Research Center; At NASA Ames Research Center, the IPG team is lead by Tony Lisotta, with supervision from Bill Johnston, Arsi Vaziri, and Tom Hinke, and support from

team leads Mary Hultquist, Warren Smith, and George Myers; plus staff members and collaborators.

The IPG team and the PACI team frequently cooperate to develop new capabilities for grids, and both teams help the Globus team with new solutions to grid problems. Grid development and deployment are based on cooperation across great distances and between diverse organizational groups.

REAL WORK EXAMPLE

Here's just one example of how grids foster collaboration among these diverse groups. Recently, a NASA research scientist, Tiffany Moisan, NASA Goddard/Wallops Flight Facility, Wallops Island, Virginia, collected microbiological samples in the tidewaters around Wallops Island, off the coast of Virginia. To see the samples at the level her research required, she needed to use a high performance microscope located at the National Center for Microscopy and Imaging Research (NCMIR), University of California, San Diego (UCSD). She sent the samples to San Diego, then used NPACI's Telescience Grid (<http://www.npaci.edu/envision/v16.2/telescience.html>) and NASA's IPG to view and control the output of the microscope from her desk on Wallops Island.

In addition to viewing the samples through the high performance microscope; she could actually move the platform holding the samples—located across the continent—and manipulate adjustments to the microscope, from Wallops Island. The microscope produces huge sets of image data; in this case the image data was stored using a Storage Resource Broker (SRB) on NASA's IPG, and Moisan was able to run algorithms on the data while watching the results in real time.

RECENT DEVELOPMENTS

A new addition to the grid community, the Distributed Terascale Facility (DTF) Project, is being built by NSF's PACI. Research institutions NCSA, SDSC, Argonne, and Caltech will work in conjunction with IBM, Intel Corporation, and Qwest Communications, Myricom, Sun Microsystems, and Oracle Corporation. The DTF is expected to perform 11.6-trillion calculations per second and store more than 450-trillion bytes of data, with a comprehensive infrastructure called the "TeraGrid" to link computers, visualization systems and data at four sites through a 40-billion bits-per-second optical network.

The British government, through the Office of Science and Technology, and with the help of IBM, is building the National Grid for collaborative scientific research in a wide spectrum of disciplines.

HOT TOPIC

The buzz at the February 2002 Global Grid Forum (www.gridforum.org) was the Open Grid Services Architecture (OGSA) being proposed by Ian Foster, Carl Kesselman, Jeffry Nick, and Steve Tuecke (Their paper, "The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration" is available at <http://www.globus.org/research/papers.html#OGSA>). An interesting aspect of the paper is that Foster, Kesselman and Tuecke, longtime spokesmen for Globus and grids, are joined here by Jeffry Nick from IBM. Here is the abstract:

In both e-business and e-science, we often need to integrate services across distributed, heterogeneous, dynamic "virtual organizations" formed from the disparate resources within a single enterprise and/or from external resource sharing and service provider relationships. This integration can be technically challenging because of the need to achieve various qualities of service when running on top of different native platforms. We present an *Open Grid Services Architecture* that addresses these challenges. Building on concepts and technologies from the Grid and Web services communities, this architecture defines a uniform exposed service semantics (the *Grid service*); defines standard mechanisms for creating, naming, and discovering transient Grid service instances; provides location transparency and multiple protocol bindings for service instances; and supports integration with underlying native platform facilities. The Open Grid Services Architecture also

defines, in terms of Web Services Description Language (WSDL) interfaces and associated conventions, mechanisms required for creating and composing sophisticated distributed systems, including lifetime management, change management, and notification. Service bindings can support reliable invocation, authentication, authorization, and delegation, if required. Our presentation complements an earlier foundational article, "The Anatomy of the Grid," by describing how Grid mechanisms can implement a service-oriented architecture, explaining how Grid functionality can be incorporated into a Web services framework, and illustrating how our architecture can be applied within commercial computing as a basis for distributed system integration—within and across organizational domains.

[Is it possible to get a good quote from someone who attended this session?]

More information is available at zdnet.com.com/2100-1105-839265.html.

DIFFICULTIES AND OPPORTUNITIES

As grids grow, grid communities will be looking for scientists and computer engineers to help build grids or to use grids for research. The NASA Advanced Supercomputing (NAS) Division at NASA Ames Research Center is a focal point for the joint university and government creation of NASA's IPG. To discuss possible opportunities for internships or collaborative research, contact Arsi Vaziri (avaziri@mail.arc.nasa.gov), IPG Deputy Project Manager.

Grid development is not for the faint of heart. As Rob Fixmer says, in eWeek, January 11, 2002, "Grid architecture, hailed in many circles as the next great evolutionary step in computer technology, is a simple concept that becomes very complex in its implementation." Imagine the difficulty of getting hundreds of brilliant developers, each with their own idea of how things should be, to cooperate enough to make geographically dispersed resources—and organizationally diverse policies—all work together. Almost all of these geniuses work anonymously; the level of cooperation obscures any clear picture of who did what. This article has attempted to name a few of the grid pioneers, but for each person mentioned, there are 20 people making a significant anonymous contributions.

The security issues alone could have prevented grids from ever coming to fruition. But brave and cooperative people have prevailed, and many grids actually work. Still, grid computing is difficult. Trying to get something—anything—to work on a grid can cause a person to swear and throw things at the wall. Building grids is hard work, as challenging and daunting as building the Transcontinental Railroad across the United States in the 1800s. An extraordinary level of courage and cooperation is required.

In a recent discussion of the perils on the path to fully functioning computational grids, an IPG staff member was asked whether the difficulties could be overcome. She replied, "We're going for it!"